Synthesizing Ranking Functions from Bits and Pieces

Caterina Urban[1,2(✉)], Arie Gurfinkel[2], and Temesghen Kahsai[2,3]

[1] ETH Zürich, Zürich, Switzerland
caterina.urban@inf.ethz.ch
[2] Carnegie Mellon University, Pittsburgh, USA
[3] NASA Ames Research Center, Moffett Field, USA

Abstract. In this work, we present a novel approach based on recent advances in software model checking to synthesize ranking functions and prove termination (and non-termination) of imperative programs.

Our approach incrementally refines a termination argument from an under-approximation of the terminating program state. Specifically, we learn *bits* of information from terminating executions, and from these we extrapolate ranking functions over-approximating the number of loop iterations needed for termination. We combine these *pieces* into piecewise-defined, lexicographic, or multiphase ranking functions.

The proposed technique has been implemented in SeaHorn – an LLVM based verification framework – targeting C code. Preliminary experimental evaluation demonstrated its effectiveness in synthesizing ranking functions and proving termination of C programs.

1 Introduction

The traditional method for proving program *termination* and other *liveness* properties is based on the synthesis of *ranking functions*, that is, for any potentially looping computation, proving that some well-founded metric strictly decreases every time around the loop.

State-of-the-art termination provers (e.g., [5,10,16]) reduce termination to the *safety* property that no program state is repeatedly visited (and it is not covered by the current termination argument), and compose termination arguments by repeatedly invoking ranking function synthesis tools (e.g., [4,8,26]).

In this work, we present a novel approach based on recent advances in *software model checking* to synthesize ranking functions and prove termination (and non-termination) of imperative programs. The core of our approach lies on an innovative use of *safety* verification techniques to build termination arguments.

This material is based upon work funded and supported by NSF Award No. 1136008 the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. This material has been approved for public release and unlimited distribution. DM-0002915.

M. Chechik and J.-F. Raskin (Eds.): TACAS 2016, LNCS 9636, pp. 54–70, 2016.
DOI: 10.1007/978-3-662-49674-9_4

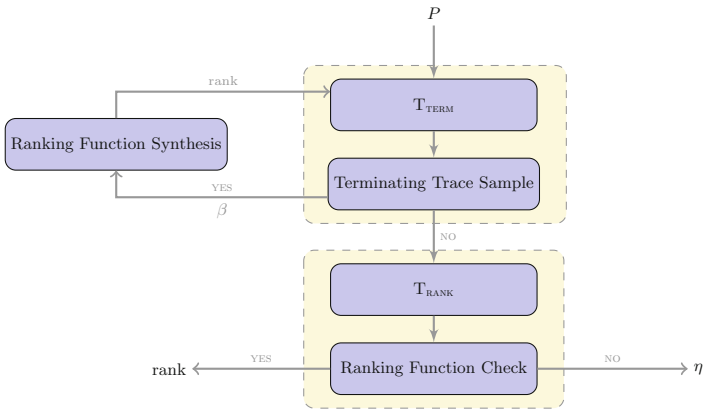

Fig. 1. Overview of our approach.

We use a safety verifier to systematically sample *terminating* program executions and extrapolate from these a candidate ranking function for the program, or to otherwise provide a witness for program non-termination. More specifically, rather than verifying that no program state is repeatedly visited, we verify the safety property that no program state is terminating (and it is not covered by the current termination argument). The counterexamples are terminating program executions which provide an *under-approximation* of the terminating program states. From these we extrapolate a candidate ranking function which *over-approximates* the number of loop iterations to termination and is possibly valid also for other terminating program executions. The candidate ranking function can be an *affine* function, or a *piecewise-defined*, *lexicographic*, or *multi-phase* combination of affine functions. We then use the safety verifier to validate that the candidate ranking function is indeed a ranking function, or to provide a counterexample non-terminating program state.

The proposed approach has been implemented in SeaHorn [15] targeting C code. We show empirically that it performs well on a wide variety of benchmarks collected from SV-COMP 2015[1], is competitive with the state-of-the-art and is able to analyze programs that are out of the reach of existing techniques.

Overview. Figure 1 provides an overview of our approach for proving termination via safety verification. The overall algorithm is presented in Sect. 3.2. A program P systematically undergoes a transformation $\mathrm{T}_{\mathrm{TERM}}$ described in Sect. 4.1 which allows sampling terminating executions β not covered by the current candidate ranking function *rank*. The candidate *rank* is systematically refined as described in Sect. 4.2 until no terminating execution β is left uncovered. Finally, P undergoes a final transformation $\mathrm{T}_{\mathrm{RANK}}$ described in Sect. 4.1 which allows validating the ranking function *rank* or providing a counterexample non-terminating state η.

[1] http://sv-comp.sosy-lab.org/2015/.

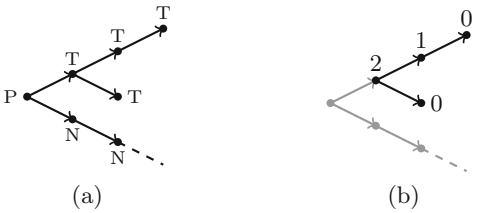

Fig. 2. Traces and ranking function.

2 Preliminaries

In this section, we introduce the basic concepts that serve in subsequent sections and we establish the notation used throughout the paper.

Transition Systems. We formalize programs using *transition systems* $\langle \Sigma, \tau \rangle$ where Σ is the set of program states and $\tau \subseteq \Sigma \times \Sigma$ defines the transition relation. Note that this model allows representing programs with (possibly unbounded) non-determinism. In the following, a program state $s \in \Sigma$ is a pair $\langle l, \bar{x} \rangle$ consisting of a program control point $l \in \mathcal{L}$ and a vector $\bar{x}$ of integers representing the values of the program variables at that control point. We write $\tau(s, s')$ for $\langle s, s' \rangle \in \tau$. The set of initial states is $\mathcal{I} \stackrel{\text{def}}{=} \{\langle i, \bar{x} \rangle \mid i \in \mathcal{L}\} \subseteq \Sigma$, where $i \in \mathcal{L}$ is the program initial control point, and the set of final states is $\mathcal{F} \stackrel{\text{def}}{=} \{\langle f, \bar{x} \rangle \mid f \in \mathcal{L}\} \subseteq \Sigma$, where $f \in \mathcal{L}$ is the program final control point.

Given a transition system $\langle \Sigma, \tau \rangle$, a *trace* is a non-empty sequence of states in Σ determined by the transition relation τ, that is $\tau(s, s')$ for each pair of consecutive states $s, s' \in \Sigma$ in the sequence. A state $s' \in \Sigma$ is *reachable* from another state $s \in \Sigma$ if and only if there exists a trace from s to s'. In the following, we write $\tau^*(s, s')$ to denote the existence of a trace from s to s'. A state $s' \in \Sigma$ is reachable if and only if it is reachable from an initial state $s \in \mathcal{I}$.

A state $s \in \Sigma$ is *terminating* if and only if all traces to which it belongs are finite, *potentially non-terminating* if and only if it belongs to at least one infinite trace. Dually, it is *non-terminating* if and only if all traces to which it belongs are infinite, and *potentially terminating* if and only if it belongs to at least one finite trace. Note that, terminating states are also potentially terminating states, and non-terminating states are also potentially non-terminating states. For instance, consider the traces depicted in Fig. 2a: the states labeled with T are terminating, the states labeled with N are non-terminating, and the state labeled with P is potentially non-terminating and potentially terminating.

Ranking Functions. The traditional method for proving termination dates back to Turing [29] and Floyd [14] and it requires finding a *ranking function*:

Definition 1 (Ranking Function). *Given a transition system* $\langle \Sigma, \tau \rangle$*, a* ranking function *is a partial function rank whose domain* dom(*rank*) *is a subset*

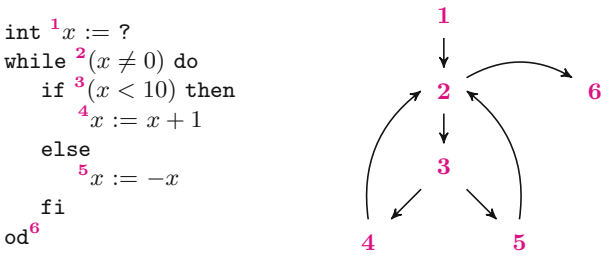

(a) (b)

Fig. 3. Terminating program 3PIECES (a) and its control flow graph (b).

of the program states and whose value (*i*) *strictly decreases through transitions between program states, that is* $\forall s, s' \in \mathrm{dom}(rank) : \tau(s, s') \Rightarrow rank(s') < rank(s)$*, and (ii) is bounded from below, that is* $\forall s \in \mathrm{dom}(rank) : rank(s) \geq 0$.

For instance, an obvious ranking function maps each program state to some well-chosen upper bound on the number of transitions until termination. Figure 2b shows a ranking function labeling the terminating states of Fig. 2a.

Control Flow Graphs. The *control flow graph* (CFG) induced by a transition system $\langle \Sigma, \tau \rangle$ is a graph whose nodes are the program control points $\mathcal{L}$ and whose edges $\mathcal{E} \subseteq \mathcal{L} \times \mathcal{L}$ are pairs of control points corresponding to transitions in the transition system: $\forall \langle l, \bar{x} \rangle, \langle l, \bar{x}' \rangle \in \Sigma : \tau(\langle l, \bar{x} \rangle, \langle l, \bar{x}' \rangle) \Rightarrow \langle l, l' \rangle \in \mathcal{E}$. In the following, we restrict our attention to *reducible* control flow graphs. A *loop* is a strongly connected component of the CFG with a single entry node h called *loop header*. The loops nested within a loop are the strongly connected components of the loop after removing the loop header. A loop *entry edge* is an edge whose source is outside the loop and whose target is inside the loop, a *loop edge* is an edge whose source and target are within the loop, and a loop *exit edge* is an edge whose source is inside the loop and whose target is outside the loop. Similarly, we can partition the corresponding transitions in the transition system into loop *entry transitions*, *loop transitions*, and loop *exit transition*.

Example 1. Consider the program in Fig. 3a: the integer variable x is initialized non-deterministically; then, at each loop iteration, the value of x is increased by one or negated when it becomes greater than or equal to ten, until x becomes zero. The control flow graph of the program is depicted in Fig. 3b. The program `while` loop corresponds to the strongly connected component of the CFG formed by the nodes **2**, **3**, **4** and **5**. The loop header is the node **2**. There is a single entry edge $\langle \mathbf{1}, \mathbf{2} \rangle$ and a single exit edge $\langle \mathbf{2}, \mathbf{6} \rangle$.

Remark 1. Note that it is not necessary for a ranking function to strictly decrease at each transition but only around each loop iteration [11]: $\forall \langle h, \bar{x} \rangle, \langle h, \bar{x}' \rangle \in \mathrm{dom}(rank) : \tau^*(\langle h, \bar{x} \rangle, \langle h, \bar{x}' \rangle) \Rightarrow rank(\langle h, \bar{x}' \rangle) < rank(\langle h, \bar{x} \rangle)$.

Example 2. The program 3PIECES of Fig. 3a terminates whatever the initial value of the variable x. The following piecewise-defined function:

$$f(x) = \begin{cases} -x & x \leq 0 \\ 21 - x & 0 < x < 10 \\ x + 1 & 10 \leq x \end{cases}$$

is a valid ranking function for the program, which maps the initial value of x to the number of loop iterations needed for termination.

3 Verifying Termination via Safety

In the late 1970s, Lamport suggested a classification of program properties into the classes of *safety* and *liveness* properties [20]. Safety properties represent requirements that should be continuously maintained by the program. On the other hand, liveness properties represent requirements that need not hold continuously but whose eventual or repeated realization must be guaranteed. Thus, a counterexample to a safety property is a *finite* (prefix of a) program execution, while for a liveness property a counterexample is an *infinite* execution on which an event of interest does not occur. A prominent example of a liveness property is *termination*. Instead, *non-termination* is a safety property since any terminating (and, thus, finite) program execution is a witness against non-termination.

3.1 Verifying Safety Properties

The verification of safety properties often amounts to checking the reachability of an *error* location: *a program is safe when the error location is unreachable*; otherwise, the program is unsafe. In the former case, safety provers often provide an *invariant* testifying the validity of the property. In the latter case, safety provers usually provide a *counterexample* trace violating the safety property. In the following, we propose some examples to informally illustrate how safety properties can be verified by checking the (un)-reachability of an error.

Verifying Non-Termination [6]. Consider the program in Fig. 4a: the integer variables x and y are initialized with value zero and nine, respectively; then, at each iteration, x and y are increased by one, until x becomes equal to y. Since safety provers report counterexample traces reaching an error location, in order to verify that the program is non-terminating, we turn terminating traces into counterexamples to be found. In Fig. 4b, we added an error location — defined as `assert(false)` — before the end of the program of Fig. 4a: only terminating traces would execute `assert(false)`, thus the program is non-terminating since in this case the error location is in fact unreachable.

```
int [1]x := 0, y := 9
while [2](x ≠ y) do
    [3]x := x + 1
    [4]y := y + 1
od[5]
```

(a)

```
int [1]x := 0, y := 9
while [2](x ≠ y) do
    [3]x := x + 1
    [4]y := y + 1
od
assert (false)[5]
```

(b)

Fig. 4. Non-terminating program (a) annotated with an error location (b).

```
int [1]x := ?, r := max{−x, 21 − x, x + 1}
while [2](x ≠ 0) do
    r := r − 1
    assert (r ≥ 0)
    if [3](x < 10) then [4]x := x + 1 else [5]x := −x fi
od[6]
```

Fig. 5. Program 3PIECES annotated with a ranking function.

Verifying a Ranking Function. Safety provers can also be used to verify whether a given function is a ranking function for a program. For instance, to check wether $\texttt{max}\{-x, 21 - x, x + 1\}$ is a ranking function for the program 3PIECES shown in Fig. 3a, we instrument the program as shown in Fig. 5: we add a variable r initialized with the given function $\texttt{max}\{-x, 21 - x, x + 1\}$; then, within the loop, according to Definition 1 and Remark 1 (i) we strictly decrease the value of r (i.e., we decrease r by one), and (ii) we assert that the value of r is bounded from below (i.e., we assert that r is greater than or equal to zero). Note that the counterexample traces that would violate the assertion are either (prefixes of) non-terminating traces, or (prefixes of) traces that are terminating but require a higher number of loop iterations with respect to the initial value of r. In this case, since the assertion is never violated, the given function $\texttt{max}\{-x, 21 - x, x + 1\}$ is a valid ranking function for the program 3PIECES.

3.2 Verifying Termination via Safety

In the following, we describe the overall algorithm for proving termination via safety. We detail our specific implementation choices in Sect. 4.

The overall algorithm is illustrated by Algorithm 1. We verify termination of each loop in a program, implicitly constructing a lexicographic ranking function for nested sets of loops [1]. The function ISTERMINATING takes as input a transition system $\langle \Sigma, \tau \rangle$ and returns either TRUE: R, meaning that the program is terminating and R is a ranking function, or FALSE: ρ, meaning that the program is potentially non-terminating and ρ is a counterexample potentially non-terminating initial state. Specifically, ISTERMINATING invokes the function ISLOOPTERMINATING for each loop in the program (identified by the function GETLOOPS, cf. Line 4) and maps each loop header h (cf. Line 3) to the

Algorithm 1. Program Termination

```
1: function IsTerminating(⟨Σ, τ⟩)
2:     R ← ∅
3:     for h ∈ getLoops(⟨Σ, τ⟩) do                ▷ h is a loop header in the program
4:         r : ρ ← isLoopTerminating(h, ⟨Σ, τ⟩)
5:         if r then                                          ▷ the loop is terminating
6:             R ← R[h ↦ ρ]
7:         else return false : ρ                ▷ ρ is a potentially non-terminating state
8:     return true : R                          ▷ R is a ranking function for the program
```

Algorithm 2. Loop Termination

```
1: function IsLoopTerminating(h, ⟨Σ, τ⟩)                          ▷ h is the loop header
2:     rank ← 0                                  ▷ candidate ranking function initialization
3:     B ← ∅
4:     while true do
5:         β ← getTerminatingTrace(h, ⟨Σ, τ⟩, rank)
6:         if β then                          ▷ there are terminating traces violating rank
7:             B ← B ∪ β
8:             rank ← getCandidateRankingFunction(rank, B)
9:         else                              ▷ there are no terminating traces violating rank
10:            η ← isRankingFunction(rank)
11:            if η then                          ▷ η is a potentially non-terminating state
12:                return false : η
13:            else                                ▷ rank is a ranking function for the loop
14:                return true : rank
```

returned ranking function (cf. Line 6), or returns as soon as a counterexample non-terminating state ρ is found (cf. Line 7). The function getLoops implements a standard control-flow analysis to identify (natural) loops within the CFG induced by the transition system $\langle \Sigma, \tau \rangle$. We omit its pseudocode due to space limitations. The identified program loops are analyzed in no specific order.

The function isLoopTerminating is shown in Algorithm 2. Initially, isLoopTerminating assumes that all program states within the loop are non-terminating and looks for a counterexample, that is, a terminating trace β (cf. Line 5). Then, the call to the function getCandidateRankingFunction computes a candidate ranking function *rank* for the (potentially terminating) states along this trace (cf. Line 8). The original non-termination property is weakened to only search for terminating traces violating the candidate *rank*, and the process starts over. The information provided by the collected terminating traces is used to incrementally refine the candidate *rank* with further ranking function pieces. In case no further terminating traces violating *rank* are found (cf. Line 9), the call to the function isRankingFunction checks wether all program states within the loop are terminating (cf. Line 10): if so, *rank* is a ranking function for the loop (cf. Line 14); if not, a counterexample potentially non-terminating

initial state η (that is, η belongs to at least one infinite trace) is returned (cf. Line 12). Note that isLoopTerminating might also not terminate (cf. Line 4).

4 Counterexample-Guided Ranking Function Synthesis

We now detail our implementation choices for the functions getTerminating-Trace, isRankingFunction and getCandidateRankingFunctions. We omit their pseudocode due to space limitations.

4.1 Search for Ranking Function Counterexamples

In Sect. 3.1, we have seen how to use a safety prover for verifying non-termination by turning terminating traces into counterexamples (cf. Fig. 4). In our approach, we use a similar intuition to systematically detect terminating traces violating a given candidate ranking function *rank*.

In the following, we consider a generic candidate *rank* and we introduce two program transformations $\mathrm{T}_{\mathrm{TERM}}$ and $\mathrm{T}_{\mathrm{RANK}}$ implemented by the functions getTerminatingTrace and isRankingFunction, respectively. We detail these transformations with respect to a specific candidate *rank* in Sect. 4.2.

$\mathrm{T}_{\mathrm{TERM}}$ Transformation. Let h be a loop header within a program $\langle \Sigma, \tau \rangle$ and let *rank* be a candidate ranking function for the loop. We modify the program in order to turn terminating traces violating *rank* into counterexamples to be found. Specifically, we modify Σ in order to include the value of *rank* and we add an error state $\omega \notin \Sigma$: $(\Sigma \times \mathbb{Z}) \cup \{\omega\}$. In the following, s, s', and $\langle h, \bar{x} \rangle$ denote program states in Σ. We also define the modified transition relation τ as follows:

- for each loop *entry transition* $\tau(s, \langle h, \bar{x} \rangle)$ there exists an entry transition τ^{rank} which also includes the candidate *rank*:

$$\tau^{rank}(\langle s, r \rangle, \langle \langle h, \bar{x} \rangle, r' \rangle) \Leftrightarrow \tau(s, \langle h, \bar{x} \rangle) \wedge r' = rank(\bar{x})$$

- for each *loop transition* $\tau(\langle h, \bar{x} \rangle, s)$ whose source is the loop header h there exists a loop transition $\tau^{\ominus}$ which also strictly decreases the value of *rank*:

$$\tau^{\ominus}(\langle \langle h, \bar{x} \rangle, r \rangle, \langle s, r' \rangle) \Leftrightarrow \tau(\langle h, \bar{x} \rangle, s) \wedge r' = r \ominus 1$$

- for each loop *exit transition* $\tau(s, s')$ there exists transition $\tau^{\lhd}$ to the error state ω when the candidate ranking function is negative:

$$\tau^{\lhd}(\langle s, r \rangle, \omega) \stackrel{\text{def}}{=} r \lhd 0$$

For every other transition $\tau(s, s')$ there exists a transition $\tau'(\langle s, r \rangle, \langle s', r' \rangle) \Leftrightarrow \tau(s, s') \wedge r' = r$. The counterexample traces that reach the error state are traces that are leaving the considered loop but violate the candidate *rank* since they require a higher number of loop iterations with respect to the initial value of *rank*. The function getTerminatingTrace returns any of these counterexamples.

```
int [1]x := ?, r := rank
while [2](x ≠ 0) do
    r := r − 1
    if [3](x < 10) then [4]x := x + 1 else [5]x := −x fi
od
assert (r ≥ 0)[6]
```

Fig. 6. Program 3PIECES annotated with a candidate ranking function *rank*.

Theorem 1. *Let h be a loop header of a program $\langle \Sigma, \tau \rangle$ and let $\langle \Sigma', \tau' \rangle$ be the program resulting from the* $\mathrm{T}_{\mathrm{TERM}}$ *transformation for a given candidate ranking function rank. Then, $\tau'^*(\langle\langle h, \bar{x}\rangle, rank(\bar{x})\rangle, \langle s, r\rangle) \wedge \tau(\langle s, r\rangle, \omega)$ if and only if there exist $s' \in \Sigma$ $\tau(s, s')$ and the transition is an exit transition, and $\tau^*(s, s')$ and the trace visits the loop header h strictly more than $rank(\bar{x})$ times.*

Example 3. Consider again the program 3PIECES of Fig. 3a. The transformation that we have just described intuitively corresponds to modifying 3PIECES as illustrated in Fig. 6: we add a variable r initialized with the candidate *rank* within the entry transition $\langle \mathbf{1}, \mathbf{2} \rangle$; then, within the loop transition $\langle \mathbf{2}, \mathbf{3} \rangle$, we decrease the value of r by one and, after the loop, we assert that the value of r is greater than or equal to zero. The assertion is equivalent to adding an error transition $\langle \mathbf{2}, \boldsymbol{\omega} \rangle$ when r is negative. The counterexample traces that violate the assertion are traces that leave the loop after $rank - r$ loop iterations, where r is the (negative) value of the variable r after the loop.

$\mathrm{T}_{\mathrm{RANK}}$ Transformation. Note that traces that never leave the considered loop are not counterexamples since they never reach the error state. For this reason Algorithm 2 includes a final validation of the ranking function (cf. Lines 10–14). We implement this using an analogous program transformation: we define entry transitions τ^{rank} and loop transitions $\tau^{\ominus}$ as before:

$$\tau^{rank}(\langle s, r\rangle, \langle\langle h, \bar{x}\rangle, r'\rangle) \Leftrightarrow \tau(s, \langle h, \bar{x}\rangle) \wedge r' = rank(\bar{x})$$

$$\tau^{\ominus}(\langle\langle h, \bar{x}\rangle, r\rangle, \langle s, r'\rangle) \Leftrightarrow \tau(\langle h, \bar{x}\rangle, s) \wedge r' = r \ominus 1$$

unlike before, for each *loop transition* $\tau(s, s')$ we also define a transition $\tau^{\lhd}$ to the error state ω when the candidate ranking function is negative:

$$\tau^{\lhd}(\langle s, r\rangle, \omega) \stackrel{\text{def}}{=} r \lhd 0$$

Other transitions are again defined as $\tau'(\langle s, r\rangle, \langle s', r'\rangle) \stackrel{\text{def}}{=} \tau(s, s') \wedge r' = r$. The counterexample traces that violate the assertion are necessarily (prefixes of) non-terminating traces, since the $\mathrm{T}_{\mathrm{TERM}}$ transformation has excluded all terminating traces violating the candidate ranking function. The function isRANKINGFUNCTION returns the initial state of any of these counterexamples.

Theorem 2. *Let h be a loop header of a program $\langle \Sigma, \tau \rangle$ and let $\langle \Sigma', \tau' \rangle$ be the program resulting from the* $\mathrm{T}_{\mathrm{RANK}}$ *transformation for a given candidate ranking function rank. Then, $\tau'^*(\langle\langle h, \bar{x}\rangle, rank(\bar{x})\rangle, \langle s, r\rangle) \wedge \tau(\langle s, r\rangle, \omega)$ if and only if $\tau^*(\langle h, \bar{x}\rangle, s)$ and the trace is the prefix of an infinite trace and visits the loop header h strictly more than $rank(\bar{x})$ times.*

Example 4. The transformation that we have just described intuitively corresponds to modifying the program 3PIECES of Fig. 3a as illustrated in Fig. 5 and described in Sect. 3.1.

4.2 Synthesis of Candidate Ranking Functions

The function GETCANDIDATERANKINGFUNCTION uses the terminating traces collected by GETTERMINATINGTRACE to extrapolate ranking function pieces which are combined into a candidate loop ranking function. We only consider *affine* pieces and leave the extrapolation of non-linear pieces for future work.

In Algorithm 2, the initial candidate is the constant function equal to zero (cf. Line 2). Then, the candidate ranking function is systematically updated in order to be valid for the newly discovered terminating traces, and possibly for other terminating traces not explicitly enumerated.

We extrapolate an affine ranking function piece from terminating traces mapping the initial states of these traces to the number of loop iterations needed for termination, and then finding an affine ranking function which fits these bits of information. More specifically, let $\{\langle \bar{x}_1, r_1\rangle, \langle \bar{x}_2, r_2\rangle, \dots\}$ be the set of pairs mapping the initial states $\bar{x}_1, \bar{x}_2, \dots$ of the collected terminating traces to the number $r_1, r_2, \dots$ of loop iterations needed for termination. We find a fitting affine function $\bar{m} \cdot \bar{x} + q$ of the program variables $\bar{x}$ by *linear interpolation*, that is by solving the system of equations:

$$\begin{aligned} \bar{m} \cdot \bar{x}_1 + q &= r_1 \\ \bar{m} \cdot \bar{x}_2 + q &= r_2 \\ &\vdots \end{aligned}$$

for the unknowns $\bar{m}$ and q.

Example 5. Let $\{\langle 9, 12\rangle, \langle 4, 17\rangle\}$ be the set of pairs mapping some initial states of the program 3PIECES of Fig. 3a to the number of loop iterations needed for termination: the initial state with $x = 9$ needs 12 loop iterations, and the initial state with $x = 4$ needs 17 loop iterations. Solving the system of equations:

$$\begin{aligned} m \cdot 9 + q &= 12 \\ m \cdot 4 + q &= 17 \end{aligned}$$

yields the affine function $21 - x$ of the program variable x. Note that this is a valid ranking function for all initial states with $0 < x < 10$, and not only for the given initial states with $x = 9$ and $x = 4$ (cf. Example 2).

When the system is unsatisfiable, we discard all collected states and we start over by building a new ranking function piece. The ranking function pieces are alternatively combined into *piecewise-defined*, *lexicographic*, or *multiphase* ranking functions [24]. These combinations have complementary strengths: piecewise-defined combinations are well-suited when multiple paths are present within loops (cf. Fig. 3a), lexicographic combinations are convenient for loops featuring unbounded non-determinism (cf. Fig. 7), and multiphase combinations target loops that go through a number of phases in their executions [3]. The choice of the combination is a parameter of the analysis.

Piecewise-Defined Ranking Functions. We represent piecewise-defined affine ranking functions using *max* combinations of affine ranking functions [25]:

$$\texttt{max}\{rank_1, \ldots, rank_n\}$$

where $rank_1, \ldots, rank_n$ are the affine ranking function pieces.

In the transformations $\mathrm{T}_{\mathrm{TERM}}$ and $\mathrm{T}_{\mathrm{RANK}}$ described in Sect. 4.1, the modified loop transitions $\tau^{\ominus}$ strictly decrease a *max* combination of ranking functions by strictly decreasing all its pieces:

$$\texttt{max}\{r_1, \ldots, r_n\} \ominus 1 = \texttt{max}\{r_1 - 1, \ldots, r_n - 1\}$$

In the added error transitions $\tau^{\lhd}$ a *max* combination of ranking functions is negative when all its pieces are negative:

$$\texttt{max}\{r_1, \ldots, r_n\} \lhd 0 \Leftrightarrow r_1 < 0 \wedge \cdots \wedge r_n < 0$$

Example 6. The transformations $\mathrm{T}_{\mathrm{TERM}}$ and $\mathrm{T}_{\mathrm{RANK}}$ of the program 3PIECES of Fig. 3a are shown in Figs. 5 and 6, respectively.

Lexicographic Ranking Functions. Lexicographic ranking functions are tuples:

$$(rank_1, \ldots, rank_n)$$

where $rank_1, \ldots, rank_n$ are affine ranking function pieces.

In the transformations $\mathrm{T}_{\mathrm{TERM}}$ and $\mathrm{T}_{\mathrm{RANK}}$, the modified loop transitions $\tau^{\ominus}$ strictly decrease a lexicographic ranking function resetting the less significant pieces to their initial affine expression:

$$(r_1, \ldots, r_i, r_{i+1}, \ldots, r_n) \ominus 1 = (r_1, \ldots, r_i - 1, rank_{i+1}, \ldots rank_n)$$

were $r_{i+1}, \ldots, r_n$ are negative and get reset to the initial $rank_{i+1}, \ldots, rank_n$. In the added error transitions $\tau^{\lhd}$ a lexicographic combination of ranking functions is negative when the first of its pieces is negative:

$$(r_1, \ldots, r_n) \lhd 0 \Leftrightarrow r_1 < 0$$

```
int [1]x := ?, y := ?, r := (x, y)
while [2](x > 0 ∧ y > 0) do
   if (snd(r) < 0) then r := (fst(r) − 1, y) else r := (fst(r), snd(r) − 1) fi
   assert (fst(r) ≥ 0)
   if [3](?) then [4]x := x − 1; [5]y := ? else [6]y := y − 1 fi
od[7]
```

Fig. 7. Program annotated with a lexicographic ranking function.

Example 7. Consider the program in Fig. 7: the integer variables x and y are initialized non-deterministically; then, at each iteration, either the value of y is decreased by one or the value of x is decreased by one and the value of y is reset non-deterministically, until either variable is less than or equal to zero. The program terminates whatever the initial value of x and y. Let (x, y) be a candidate lexicographic ranking function for the program. In this case, the transformation $\mathrm{T}_{\mathrm{RANK}}$ intuitively corresponds to adding a variable r initialized with (x, y) within the entry transition $\langle \mathbf{1}, \mathbf{2} \rangle$; then, within the loop transition $\langle \mathbf{2}, \mathbf{3} \rangle$, decreasing the value of r lexicographically *resetting* its second component $\mathtt{snd}(r)$ when negative, and asserting that its first component $\mathtt{fst}(r)$ is greater than or equal to zero. The assertion is equivalent to adding an error transition $\langle \mathbf{2}, \boldsymbol{\omega} \rangle$ when $\mathtt{fst}(r)$ is negative. In this case, since the assertion is never violated, (x, y) is a valid lexicographic ranking function for the program.

Multiphase Ranking Functions. Multiphase ranking functions specify ranking functions that proceed through a certain number of phases during program execution [24]. They are represented as tuples:

$$(rank_1, \ldots, rank_n)$$

where $rank_1, \ldots, rank_n$ are affine ranking function pieces. Each piece represents a phase of the ranking function. In the transformations $\mathrm{T}_{\mathrm{TERM}}$ and $\mathrm{T}_{\mathrm{RANK}}$, the modified loop transitions $\tau^{\ominus}$ strictly decrease a multiphase combination of ranking functions as follows:

$$(r_1, \ldots, r_i, r_{i+1}, \ldots, r_n) \ominus 1 = (r_1, \ldots, r_i - 1, r_{i+1}, \ldots r_n)$$

were $r_{i+1}, \ldots, r_n$ are negative (and, unlike in the lexicographic combination, are never reset). In the added error transitions $\tau^{\lhd}$ a multiphase combination of ranking functions is negative when the first of its pieces is negative:

$$(r_1, \ldots, r_n) \lhd 0 \Leftrightarrow r_1 < 0$$

In summary, our approach systematically collects terminating program executions and searches for a function that uniformly captures the termination argument of the program. The function can be an affine ranking function, or a piecewise, lexicographic, or multiphase combination of affine functions. Then, we either manage to validate the candidate ranking function or we provide a witness for program non-termination.

	Tot	Time
SeaHorn	135	1.71s
AProVE [28]	129	10.77s
FuncTion [30]	111	0.55s
HIPTnT+ [22]	152	0.62s
Ultimate [16]	109	8.45s

(a)

	SeaHorn			
	■	●	×	▲
AProVE [28]	39	33	96	22
FuncTion [30]	50	26	85	29
HIPTnT+ [22]	16	33	119	22
Ultimate [16]	55	29	80	26

(b)

Fig. 8. Overview of the experimental evaluation.

5 Implementation

Our approach is implemented in SeaHorn[2], an LLVM [21] based safety verification framework. SeaHorn verifies user-supplied assertions as well as a number of built-in safety properties (e.g., buffer and signed integer overflows). It can also be used to check for inconsistent code in C programs [18].

SeaHorn is parameterized by the semantic representation of the program using Constrained Horn Clauses (CHCs), and by the verification engine that leverages the latest advances made in SMT-based Model Checking and Abstract Interpretation. Detailed information about SeaHorn can be found in [15]. The transformations T_{TERM} and T_{RANK} presented in Sect. 4.1 are used to enhance the CHCs passed to the verification engine. SeaHorn employs several SMT-based model checking engines based on PDR/IC3 [2], including Spacer [19]. The synthesis of candidate ranking functions presented in Sect. 4.2 uses Z3 [12] to find affine functions fitting the collected terminating states.

Experimental Evaluation. We compared SeaHorn to the participants in the termination division of SV-COMP 2015: AProVE [28], FuncTion [30], HIPTnT+ [22], and Ultimate Automizer [16]. We evaluated the tools against 190 terminating C programs collected from the SV-COMP 2015 benchmarks. Specifically, we selected only the programs that *all* tools could analyze (e.g., without parse errors or other clear issues) among the two most populated verification tasks of the termination category (i.e., crafted-lit and memory alloca). Note that other tools (e.g., FuncTion) provide a very limited support for arrays and pointers. Therefore, we were not able to analyze 30 % of the considered benchmarks. The experiments were performed on a machine with a 2.90 GHz 64-bit Dual-Core CPU (Intel i5-5287U) and 4 GB of RAM, and running Ubuntu 14.04.

In the evaluation, we run in parallel three instances of SeaHorn parameterized with the different ranking function combinations presented in Sect. 4.2, halting the analysis as soon as one instance reported a result. Figure 8 summarizes our experimental evaluation and Fig. 9 shows a detailed comparison of SeaHorn against each other tool. In Fig. 8a, the first column reports the total

[2] http://seahorn.github.io/.

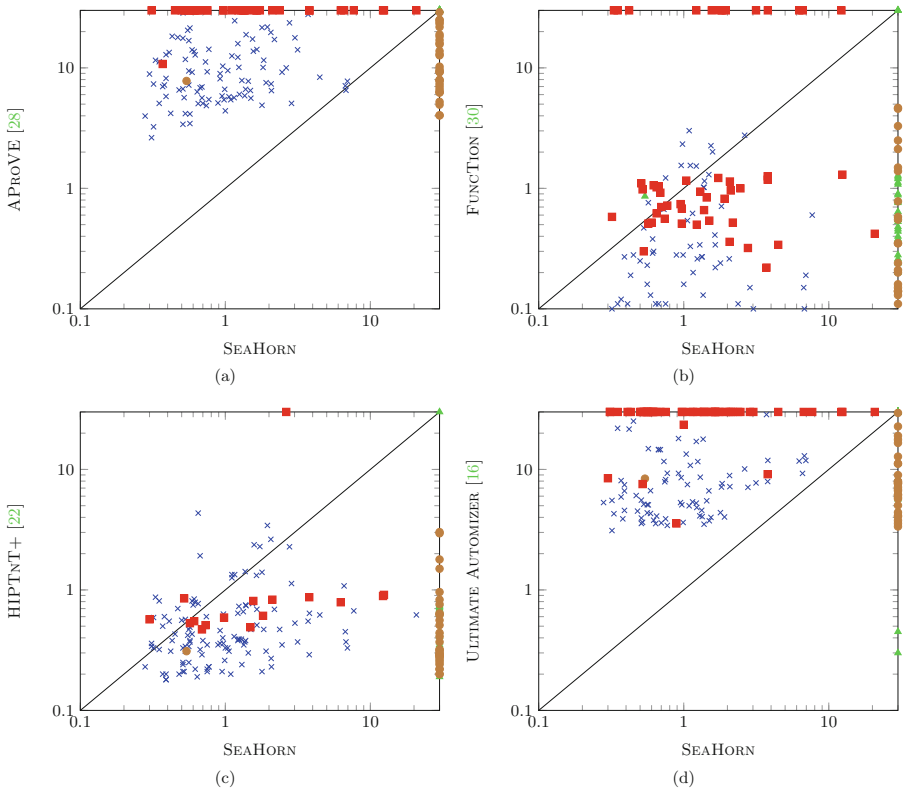

Fig. 9. Detailed comparison of SeaHorn against AProVE [28] (a), FuncTion [30] (b), HIPTnT+ [22] (c), and Ultimate Automizer [16] (d).

number of programs that each tool could prove terminating, and the second column reports the average running time in seconds for the programs where the tool proved termination. We used a time limit of 30 s for each program. In Fig. 8b, the first column (■) lists the total number of programs that the tool was not able to prove termination for and that SeaHorn could prove terminating, the second column (●) reports the total number of programs that SeaHorn was not able to prove termination for and that the tool could prove terminating, and the last two columns report the total number of programs that both the tool and SeaHorn were able (×) or unable (▲) to prove terminating. The same symbols are used in Fig. 9.

Figure 8a shows that SeaHorn is able to prove termination of 3.2 % more programs than AProVE, 12.6 % more programs than FuncTion, and 13.7 % more programs than Ultimate Automizer. HIPTnT+ is able to prove termination of 8.9 % more programs than SeaHorn, but SeaHorn can prove termination of 42.1 % of the programs that HIPTnT+ is not able to prove terminating (8.4 % of the total program test cases, cf. Fig. 8b).

Figure 8b highlights the complementary strengths of SeaHorn and each of the other tools. Specifically, SeaHorn and AProVE seem to form the best combination with respectively 20.5 % and 17.4 % of the total program test cases that could be proved terminating only by one tool and not the other, and only 11.6 % of the test case that could not be proved terminating by either tool.

Figure 9 shows that SeaHorn is generally faster than AProVE (cf. Fig. 9a) and Ultimate Automizer (cf. Fig. 9d), and often slower than FuncTion (cf. Fig. 9b) and HIPTnT+ (cf. Fig. 9c). In Fig. 9b and c, we also see that FuncTion and HIPTnT+ give up earlier when unable to prove termination, while SeaHorn, AProVE, and Ultimate Automizer usually persist with the analysis until the timeout (cf. also Fig. 9a and d).

Finally, we noticed that five of the *SV-COMP 2015* program test cases could be proved terminating only by SeaHorn (one only by AProVE, one only by FuncTion, two only by HIPTnT+, and five only by Ultimate). No tool could prove termination of six of the program test cases.

6 Related Work

In the recent past, termination analysis has benefited from many research advances and powerful termination provers have emerged. Many approaches in this area reduce termination to a safety property. For instance, the approach implemented in Terminator [10] systematically verifies that no program state is repeatedly visited (and it is not covered by the current termination argument). The identified counterexamples are independently proved to be terminating [26] building a disjunctive well-founded termination argument [27]. A similar incremental approach is used in T2 [5] for the construction of lexicographic ranking functions. An automata-based incremental approach is described in [17] and implemented in Ultimate [16]. An approach based on conflict-driven learning is used in [13] to enhance the abstract interpretation-based termination analysis [31] implemented in FuncTion [30].

The incremental approach that we have proposed in this paper uses safety verifiers for proving termination in a fundamentally different way than existing methods: rather than systematically verifying that no program state is visited repeatedly, we systematically verify that no program state is terminating. Thus, our counterexamples are finite traces and do not need to be proven terminating.

The counterexample finite traces identified by our approach are used to extrapolate affine ranking functions. The linear interpolation that we use resembles the widening operator described in [31]. The extrapolated ranking functions are combined into a piecewise-defined, lexicographic, or multiphase ranking function for a program. Thus, our method provides more valuable information than just a positive or inconclusive answer like the methods based on the size-change termination principle [23] and implemented in AProVE [28], or like the already cited methods based on disjunctive well-foundedness and implemented in Terminator. Finally, compared to the incomplete methods implemented in AProVE and FuncTion, our method is also able to prove program non-termination.

7 Conclusion and Future Work

This paper provides a new perspective on the use of safety verifiers for proving program (non-)termination. We have proposed a novel incremental approach, which uses a safety verifier to systematically sample *terminating* program executions and synthesize from these a ranking function for the program, or to otherwise provide a witness for program non-termination.

It remains for future work to adapt the approach in order to infer sufficient preconditions for program termination [7,31]. We also plan to extend the approach to other liveness properties [9,32].

References

1. Ben-Amram, A.M.: Ranking functions for linear-constraint loops. In: VPT, pp. 1–8 (2013)
2. Bradley, A.R.: IC3 and beyond: incremental, inductive verification. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, p. 4. Springer, Heidelberg (2012)
3. Bradley, A.R., Manna, Z., Sipma, H.B.: The polyranking principle. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 1349–1361. Springer, Heidelberg (2005)
4. Bradley, A.R., Manna, Z., Sipma, H.B.: Termination analysis of integer linear loops. In: Abadi, M., de Alfaro, L. (eds.) CONCUR 2005. LNCS, vol. 3653, pp. 488–502. Springer, Heidelberg (2005)
5. Brockschmidt, M., Cook, B., Fuhs, C.: Better termination proving through cooperation. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 413–429. Springer, Heidelberg (2013)
6. Chen, H.-Y., Cook, B., Fuhs, C., Nimkar, K., O'Hearn, P.W.: Proving nontermination via safety. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014. LNCS, vol. 8413, pp. 156–171. Springer, Heidelberg (2014)
7. Chen, H.Y., David, C., Kroening, D., Schrammel, P., Wachter, B.: Synthesising interprocedural bit-precise termination proofs. In: ASE (2015)
8. Colón, M.A., Sipma, H.B.: Synthesis of linear ranking functions. In: Margaria, T., Yi, W. (eds.) TACAS 2001. LNCS, vol. 2031, pp. 67–81. Springer, Heidelberg (2001)
9. Cook, B., Khlaaf, H., Piterman, N.: On automation of CTL* verification for infinite-state systems. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015, Part I. LNCS, vol. 9206, pp. 13–29. Springer, Heidelberg (2015)
10. Cook, B., Podelski, A., Rybalchenko, A.: Termination proofs for systems code. In: PLDI, pp. 415–426 (2006)
11. Cousot, P., Cousot, R.: An abstract interpretation framework for termination. In: POPL, pp. 245–258 (2012)
12. de Moura, L., Bjørner, N.S.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
13. D'Silva, V., Urban, C.: Conflict-driven conditional termination. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015, Part II. LNCS, vol. 9207, pp. 271–286. Springer, Heidelberg (2015)

14. Floyd, R.W.: Assigning meanings to programs. Proc. Symp. Appl. Math. **19**, 19–32 (1967)
15. Gurfinkel, A., Kahsai, T., Komuravelli, A., Navas, J.A.: The seahorn verification framework. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015, Part I. LNCS, vol. 9206, pp. 343–361. Springer, Heidelberg (2015)
16. Heizmann, M., Dietsch, D., Leike, J., Musa, B., Podelski, A.: Ultimate automizer with array interpolation (competition contribution). In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 455–457. Springer, Heidelberg (2015)
17. Heizmann, M., Hoenicke, J., Podelski, A.: Software model checking for people who love automata. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 36–52. Springer, Heidelberg (2013)
18. Kahsai, T., Navas, J.A., Jovanovic, D., Schäf, M.: Finding inconsistencies in programs with loops. In: Davis, M., et al. (eds.) LPAR-20 2015. LNCS, vol. 9450, pp. 499–514. Springer, Heidelberg (2015). doi:10.1007/978-3-662-48899-7_35
19. Komuravelli, A., Gurfinkel, A., Chaki, S.: SMT-based model checking for recursive programs. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 17–34. Springer, Heidelberg (2014)
20. Lamport, L.: Proving the correctness of multiprocess programs. IEEE Trans. Softw. Eng. **3**(2), 125–143 (1977)
21. Lattner, C., Adve, V.S.: LLVM: a compilation framework for lifelong program analysis & transformation. In: CGO, pp. 75–88 (2004)
22. Le, T.-C., Qin, S., Chin, W.-N.: Termination and non-termination specification inference. In: PLDI, pp. 489–498 (2015)
23. Lee, C.S., Jones, N.D., Ben-Amram, A.M.: The size-change principle for program termination. In: POPL, pp. 81–92 (2001)
24. Leike, J., Heizmann, M.: Ranking templates for linear loops. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014. LNCS, vol. 8413, pp. 172–186. Springer, Heidelberg (2014)
25. Ovchinnikov, S.: Max-min representation of piecewise linear functions. Contrib. Algebra Geom. **42**(1), 297–302 (2002)
26. Podelski, A., Rybalchenko, A.: A complete method for the synthesis of linear ranking functions. In: Steffen, B., Levi, G. (eds.) VMCAI 2004. LNCS, vol. 2937, pp. 239–251. Springer, Heidelberg (2004)
27. Podelski, A., Rybalchenko, A.: Transition invariants. In: LICS, pp. 32–41 (2004)
28. Ströder, T., Aschermann, C., Frohn, F., Hensel, J., Giesl, J.: AProVE: termination and memory safety of C programs (competition contribution). In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 417–419. Springer, Heidelberg (2015)
29. Turing, A.: Checking a large routine. In: Report of a Conference on High Speed Automatic Calculating Machines, pp. 67–69 (1948)
30. Urban, C.: FuncTion: an abstract domain functor for termination (competition contribution). In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 464–466. Springer, Heidelberg (2015)
31. Urban, C., Miné, A.: A decision tree abstract domain for proving conditional termination. In: Müller-Olm, M., Seidl, H. (eds.) SAS 2014. LNCS, vol. 8723, pp. 302–318. Springer, Heidelberg (2014)
32. Urban, C., Miné, A.: Proving guarantee and recurrence temporal properties by abstract interpretation. In: D'Souza, D., Lal, A., Larsen, K.G. (eds.) VMCAI 2015. LNCS, vol. 8931, pp. 190–208. Springer, Heidelberg (2015)